

WO02082709

Publication Title:

High-speed message forwarding lookups for arbitrary length strings using pipelined memories

Abstract:

The invention provides a method and system for lookup of message header information that has the advantages of low-power, speed, and flexibility. The invention uses a sequence of pipelined on-chip memories, each having only a portion of the header information intended for lookup. Each one of the on-chip memories simultaneously performs a lookup on a portion of the header information, allowing embodiments of the invention to operate on multiple messages worth of header information substantially simultaneously. The invention uses a novel data structure for recording destination addresses in the sequence of on-chip memories, having the property that moving information about destination addresses, or otherwise responding to changes in network topology is flexible, while at the same time maintaining relatively dense usage of the on-chip memories. The novel data structure allows embodiments of the invention to find the longest match destination address with no substantial additional processing over random-access memory lookup

Data supplied from the esp@cenet database - <http://ep.espacenet.com>

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
17 October 2002 (17.10.2002)

PCT

(10) International Publication Number
WO 02/082709 A2

(51) International Patent Classification⁷: H04L
(21) International Application Number: PCT/US02/10796
(22) International Filing Date: 4 April 2002 (04.04.2002)
(25) Filing Language: English
(26) Publication Language: English
(30) Priority Data:
09/827,270 4 April 2001 (04.04.2001) US
(71) Applicants: CYPRESS SEMICONDUCTOR CORPORATION [US/US]; 3901 North First Street, San Jose, CA 95134 (US). VENKATACHARY, Srinivasan [—/US]; 1064 Reed Avenue, Apt. #10, Sunnyvale, CA 94086 (US).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZM, ZW.

(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:


— without international search report and to be republished upon receipt of that report

(72) Inventor: GUPTA, Pankaj; 260 Lassen Avenue, Mountain View, CA 94043 (US).

(74) Agent: SWERNOFSKY, Steven, A.; Swernofsky Law Group PC, P.O. Box 390013, Mountain View, CA 94039-0013 (US).

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: HIGH-SPEED MESSAGE FORWARDING LOOKUPS FOR ARBITRARY LENGTH STRINGS USING PIPELINED MEMORIES



(57) Abstract: The invention provides a method and system for lookup of message header information that has the advantages of low-power, speed, and flexibility. The invention uses a sequence of pipelined on-chip memories, each having on a portion of the header information intended for lookup. Each one of the on-chip memories simultaneously performs a lookup on a portion of the header information, allowing embodiments of the invention to operate on multiple messages worth of header information substantially simultaneously. The invention uses a novel data structure for recording destination addresses in the sequence of on-chip memories, having the property that moving information about destination addresses, or otherwise responding to changes in network topology is flexible, while at the same time maintaining a relatively dense usage of the on-chip memories. The novel data structure allows embodiments of the invention to find the longest match destination address with no substantial additional processing over random-access memory lookup.

HIGH-SPEED MESSAGE FORWARDING LOOKUPS FOR ARBITRARY LENGTH STRINGS USING PIPELINED MEMORIES

Background of the Invention

5

1. Field of the Invention

This invention relates to message forwarding, such as for example in a computer network, wireless network or intra or inter computer communication
10 system.

2. Related Art

In interconnection networks, individual devices communicate with each
15 other using messages, which are forwarded using a collection of routers (or switches) from a source device to a destination device. Messages include header information, which is used by the routers to determine how to treat those messages. For example, each router must decide which next-hop router to forward a message to so that the message ultimately reaches its proper destination.

20

There are at least three concerns each router must address when forwarding messages. First, routers perform a lookup for routing information in response to packet header information. Second, routers schedule the packet for sending in response to a result of the lookup operation. Third, routers forward the
25 packet to an output queue for sending, and actually send the packet to a next hop location.

When performing the lookup operation, there are further concerns each router must address. First, routers need a rapid technique for determining how to treat
30 the message in response to the destination address, including obtaining the "longest prefix match" for the header information. Second, routers need an efficient technique

for recording all known destination prefixes addresses in memory and for updating that memory in response to changes in network topology. (While this operation does not take place for each lookup, the router still performs the operation from time to time.) Given the so-far rapid growth of the Internet, and the consequent rapid growth
5 in both the number of possible destinations and the number of routing table entries, every step in this process requires accuracy and would benefit from speed.

The “longest prefix match” problem may be shown by an example. The destination address of a message directed to MIT maybe routed a first way; such
10 messages may include a prefix that is broad and general. The destination address of a message directed to MIT’s School of Architecture may be routed a second way, different from the first way and having priority over the first way. These distinctions and priorities are made by looking up the destination address in a table that associates a destination address with a next hop router and give preference to table entries that
15 match the longest part of the destination address (this is referred to as “longest match”). Given the rapid expansion of the Internet, tables are relatively large and growing relatively rapidly. Hence, identifying the “longest match” has become increasingly difficult.

20 Longest prefix matching can be performed using a TCAM (ternary content addressable memory) or by using various different algorithms such as different trie structures. Both TCAMs and the various different algorithms have numerous drawbacks. TCAMS have relatively low memory density, perform at a limited clock speed, and so require multiple chips in order to achieve either relatively
25 high speeds or relatively large table sizes.

Moreover, algorithms used in lookups have relatively large and elaborate data structures, thus requiring the use of off-chip memory (additional memory not integrated onto the same chip as the lookup circuits).

Known methods of matching destination addresses primarily use off-chip memory, in which the router (or switch) communicates with a memory device using a memory bus to retrieve information about routing messages. When off-chip memory is used, the memory can be either very dense but relatively slow, such as DRAM, or can be very fast but relatively less dense, such as off-chip SRAM. Using DRAM, the speed with which the router can reference memory (and the amount of bandwidth for memory access) limits the speed at which the router can process header information. The relative lesser density of off-chip SRAM means that the router uses multiple SRAM devices for the same amount of storage, and therefore devotes a relatively large number of its input/output pins to accessing memory (also limiting the ability to quickly access large amounts of memory). Either of these drawbacks severely limits the speed with which the router can operate.

Alternatively, routers may use “on-chip” memory, in which the router includes memory with routing information integrated onto the same monolithic semiconductor circuit with intelligent logic. However, on-chip SRAM memory sizes are relatively limited, so that relatively compact or small data structures are preferred. On-chip DRAM memory sizes are less limited, but on-chip DRAM is comparatively slow, and therefore less suitable for rapid lookup.

TCAMs have the advantage of not requiring large and elaborate data structures, because they are specially designed for the purpose of lookup, but they have relatively high power usage, low speed and low density. Avoiding specialized memories leads to using other data structures, which in the known art are relatively large and therefore use multiple memory devices—either multiple SRAM chips (to achieve larger storage), or multiple DRAM chips (to achieve higher speed).

Accordingly, it would be desirable to provide a technique for lookup of message header information that is not subject to drawbacks of the known art. This can be achieved using aspects of the invention in which a relatively compact and smaller data structure is used in combination with pipelined on-chip memory.

Summary of the Invention

The invention provides a method and system for lookup of message header information that has the advantages of (1) the relative low scaling factor for power usage associated with the use of a random access memory such as SRAM, which has relatively constant power requirement relative to the amount of stored data, over that of associative memories such as TCAM, which has power requirement proportional to the amount of stored data, (2) the relative flexibility, associated with the use of a random access memory such as SRAM, over that of more specialized memories such as TCAM, (2) the relative speed associated with on-chip memory over that of off-chip memory or associative memory. The invention includes methods and systems for lookup using pipelined on-chip memory to access any relatively compact and smaller data structures stored on-chip; one such relatively compact data structure is described in the incorporated disclosure, as referred to herein.

In a first aspect of the invention, lookups are performed using a sequence of pipelined on-chip memories, each having only a portion of the header information intended for lookup. Each one of the sequence of on-chip memories simultaneously performs a lookup on a portion of the header information, thus allowing embodiments of the invention to operate on multiple messages worth of header information substantially simultaneously. Even though each one lookup might take the full sequence of on-chip memories to complete, the router is still able to receive and to process one message for each memory cycle. Lookup speed is limited only by on-chip memory speed, rather than by additional limitations of the known art. Given that on-chip memory is the fastest of all available types of memory, lookup speed (and therefore router capability) scales well with raw technology. Moreover, the use of a pipelined on-chip memory is particularly cost effective.

The invention has general applicability to compact memory storage and retrieval, to update of information recorded in relatively compact form, and to applications having utility for data lookup. None of these applications are limited

specifically to lookup for message header information, nor are they necessarily related to the specific applications disclosed herein. For example, embodiments of the invention can include one or more of, or some combination of, the following applications:

5

- Routers and switches at all levels of the ISO/OSI networking model, including without limitation: bridges, network gateways, network routers and switches.

10

- Emulation of an associative memory with contiguous prefix masking, including applications other than routing or switching.

15

- Both fixed-length and variable-length lookups in various types of lookup tables, such as for example for different protocols such as IP, Ethernet, MPLS, and different message types, such as for example unicast, multicast, and the like.

Brief Description of the Drawings

Figure 1 shows a block diagram of a system for pipelined memory organization for flexible prefix matching.

20

Figure 2 shows a process flow diagram of a method for storage and retrieval of information using a system for pipelined memory organization for flexible prefix matching.

25

Description of the Preferred Embodiment

The invention is described herein with regard to preferred steps and data structures. Those skilled in the art will recognize, after perusal of this application, that the described steps and data structures are not limited to any particular circuits. Rather, those of ordinary skill in the art would be able to

30

implement the described steps and data structures, and equivalents thereof, without undue experimentation or further invention. All such implementations are within the scope and spirit of the invention.

5 *Related Application*

Inventions described herein can be used in conjunction with inventions described in the following document:

- 10 • U.S. Patent Application Serial No. 09/826,556, filed 4 April 2001 in the name of Pankaj Gupta and Srinivasan Venkatachary, titled "Compact Data Structures for Pipelined Message Forwarding Lookups."

This document is hereby incorporated by reference as if fully set forth
15 herein. This document is referred to as the incorporated disclosure.

A preferred embodiment includes data structures and memories as further described in the incorporated disclosure.

20 *Lexicography*

The following terms refer or relate to aspects of the invention as described below. The descriptions of general meanings of these terms are not intended to be limiting, only illustrative.

25

- **router** — in general, any device for performing lookup of message header information against an information database, and for treating messages having that message header information differentially in response to a result of the lookup. For example, a router can act in response to destination IP addresses (for forwarding), in response to destination IP addresses and source IP
30 addresses (for multicast forwarding and for access control), or in response to

other packet header fields (for enforcing administrative restrictions or other message routing rules).

As noted above, these descriptions of general meanings of these terms are not intended to be limiting, only illustrative. Other and further applications of the invention, including extensions of these terms and concepts, would be clear to those of ordinary skill in the art after perusing this application. These other and further applications are part of the scope and spirit of the invention, and would be clear to those of ordinary skill in the art, without further invention or undue experimentation.

System Elements

Figure 1 shows a block diagram of a system for pipelined memory organization for flexible prefix matching.

A system 100 includes a lookup circuit 101, and can include an (optional) off-chip memory 105, coupled to the lookup circuit 101 using a bus 106 (such as for example a memory bus).

As noted above, the off-chip memory 105 and the bus 106 are optional. A preferred embodiment does not require them, but can include them if desired to provide additional memory accessible by the lookup circuit 101.

In a preferred embodiment, the lookup circuit 101 includes a single monolithic circuit integrated onto at least one side of a silicon wafer ("chip"). Thus, the lookup circuit 101 is preferably integrated onto a single chip, and the sequence of SRAM memories (as further described below) is thus accessible by techniques for accessing "on-chip" memory.

The lookup circuit 101 includes an input port 102, an (optional) pre-lookup circuit 103, a sequence of SRAM memories (as further described below), and an output port 104.

5 There is no particular requirement that the sequence of memories must be SRAM; they may be DRAM or may other type of on-chip memory or on-chip storage. References to SRAM memory herein should be broadly read to encompass any method of on-chip storage, now known or discovered in the future.

10 In a preferred embodiment, the input port 102 can be coupled to the bus 106, so as to be included in a larger system in which packet header information is sent to the lookup circuit 101. However, in alternative embodiments, the input port 102 can be directly coupled to an output of a circuit that delivers packet header information, or configured to receive packet header information (or other
15 information) in some other way.

 The (optional) pre-lookup circuit 103 includes a set of pre-lookup logic 115 and a pre-lookup memory circuit 120. Information coupled to the input port 102 is coupled to the pre-lookup logic 115 and the pre-lookup memory circuit 120. The
20 pre-lookup memory circuit 120, under control of the pre-lookup logic 115, performs a direct memory lookup for prefix values that are relatively short (that is, smaller than a value d_0 , a selected number of bits). In embodiments of the invention, d_0 can be any value, so long as that value is sufficiently small to allow the lookup circuit 101 to operate effectively (for example, the pre-lookup memory circuit 120 must be fit onto
25 the chip, and there may be other constraints). In a preferred embodiment, $d_0 = 14$, so prefix values of 14 or fewer bits are all recorded for lookup by the pre-lookup memory circuit 120, which can include a 2^{14} -entry table.

 The packet header information, having a total bit width of w , is coupled
30 to the sequence of SRAM memories. Each stage n in the sequence includes an on-chip SRAM memory, an address lookup circuit, and a pipeline register, and spans d_n

bits of the total bit width. Thus, a first stage in the sequence includes a first on-chip SRAM memory 125, a first address lookup circuit 130, and a first pipeline register 135, and spans d_1 bits of the total bit width; a second stage in the sequence includes a second on-chip SRAM memory 140, a second address lookup circuit 145, and a second pipeline register 150, and spans d_2 bits of the total bit width; and so on until a final stage in the sequence includes a final on-chip SRAM memory and a final address lookup circuit 160 (except as described below, there is no particular requirement for a final pipeline register).

10 In a preferred embodiment, each value d_n is selected for optimal processing of a lookup search key. In alternative embodiments, each value d_n may be distinct or identical, may be selected to optimize a particular type of lookup, or may be arbitrarily chosen.

15 In a preferred embodiment, while each SRAM memory has only a selected number of entries, those entries can be quite wide in terms of the number of bits of information they provide, within only the limits of integrating SRAM memories onto a chip. Thus for example, the result provided by an individual entry within one of the SRAM memories can be a data structure having 256 bits, 1K bits, or
20 more, depending on the nature of the data stored therein and the physical size of the SRAM memories as integrated onto the chip.

A result of the lookup operation is coupled to the corresponding stage n address lookup circuit for interpretation. As described below, the lookup operation
25 can provide a final lookup result (or a portion thereof) when there is a prefix entry for a prefix length within the stride of d_n bits within the range of the corresponding stage n SRAM memory. Also as described below, the lookup operation can provide a further lookup index into the corresponding next stage $n+1$ SRAM memory. The address lookup circuit determines from the stage n lookup result whether it includes a
30 final lookup result or a further lookup index.

There is no particular requirement that the further lookup index must point into the corresponding next stage $n+1$ SRAM memory. In some data structures, the further lookup index may point into a later stage on-chip memory (a “jumping pointer” forward into the sequence of memories), or may point into an earlier stage on-chip memory (a “backward pointer”, particularly for loopback or multiple-loopback operation).

An output of the address lookup circuit is coupled to the corresponding stage n pipeline register; the pipeline registers are clocked (preferably using a common clock) so that the lookup operation can proceed at each stage n substantially in parallel. Thus, as described herein, lookup #1 is being partially performed at stage n , lookup #2 is being partially performed at stage $n-1$, ..., lookup # $n-1$ is being performed at stage 2, and lookup # n is being performed at stage 1, all substantially simultaneously.

When a lookup is completed, the final lookup result is coupled to the output port 104.

Loopback Functionality

In a first set of alternative embodiments, the lookup search key may be longer than the width of the sequence of memories, and any remaining portion of the lookup search key not yet referenced may be looped back to an earlier part (such as for example the beginning) of the sequence. Thus, for example, where the packet header information has a total bit width of $2w$, the second half can be looped back to the beginning of the sequence of memories, for further processing.

Loopback can be performed multiple times. Where the packet header information has a total bit width of $3w$, the second third can be looped back to the beginning of the sequence of memories for further processing, and the final third can

be looped back a second time to the beginning of the sequence of memories for further processing.

In this first set of alternative embodiments, a result for the first w bits of the lookup search key is maintained in the final pipeline register, and forwarded with the portion of the lookup search key not yet processed (the second w bits in those cases when the lookup search key is $2w$ bits wide, or the remaining $2w$ bits in those cases when the lookup search key is $3w$ bits wide, and the like) to the stage 1 memory, or to a later stage memory when the lookup search key is not an integer multiple of w bits.

Loopback need not be the very next lookup operation performed. In a preferred embodiment, any intermediate state is recorded with the loopback along with the remaining part of the lookup search key, and the entire looped-back lookup operation is queued with other lookup operations or update operations for being performed in a priority order. This prevents looped-back lookup operations from taking too much of the memory lookup resources available.

This first set of alternative embodiments may be combined with the second set of alternative embodiments, as described below. Thus, the lookup search key may be any width; if more than w bits, the loopback functionality may be used to process the portion of the lookup search key exceeding w bits, while if less than w bits, the multiple-issue functionality (as described below) may be used to insert the lookup search key into a stage n memory for n other than 1.

Multiple-Issue Functionality

In a second set of alternative embodiments, the lookup search key may be inserted into the sequence of memories, and the lookup result may be provided from the sequence of memories, at points other than the beginning and end of the sequence. This has the advantages of (1) allowing for early reporting of lookup

results where appropriate; (2) allowing for parallel processing of multiple lookup search keys at once. Thus, for example, if there are 24 stages in the sequence of memories, the first 8 stages may be used for a first lookup, the second 8 stages may be used for a second lookup, and the third 8 stages may be used for a third lookup, all
5 in parallel, thus providing a tripled lookup rate.

In this second set of alternative embodiments, a set of multiple-issue logic (not shown) inserts the lookup search key into the pipeline of memories at a stage n memory (for n other than 1) and extracts the lookup result from the pipeline
10 of memories at a stage n memory (for n other than w). The multiple-issue logic determines, for each lookup search key, where to insert the lookup search key and where to extract the lookup result, so that parallel multiple issue of lookup results may be achieved for multiple lookup search keys. In the example above, the multiple-issue logic would determine which subsequence of the pipelined memories each
15 lookup search key would be assigned to.

As noted above, the first set of alternative embodiments may be combined with the second set of alternative embodiments, so that portions of lookup search keys less than w bits may be processed using multiple-issue functionality while
20 portions of lookup search keys when those lookup search keys exceed w bits may be processed using loopback functionality.

Method of Use

25 Figure 2 shows a process flow diagram of a method for storage and retrieval of information using a system for pipelined memory organization for flexible prefix matching.

A method 200 includes a set of flow points and process steps as
30 described herein.

Although by the nature of textual description, the flow points and process steps are described sequentially, there is no particular requirement that the flow points or process steps must be sequential. Rather, in preferred embodiments of the invention, the described flow points and process steps are performed in a parallel or pipelined manner.

At a flow point 210, the lookup circuit 101 is ready to receive a lookup search key at the input port 102.

At a step 211, the lookup circuit 101 couples the lookup search key to the pre-lookup circuit 103.

In a preferred embodiment, the lookup circuit 101 also couples an instruction, such as a set of flag bits (not shown) indicating whether the lookup operation is intended to store information in, or retrieve information from, the pipelined memory organization. Although this functionality is described herein as “storing” information, in a preferred embodiment the flag bits can indicate any kind of modify or update operation, such as to clear (that is, delete) an entry, mark an entry invalid, modify an entry to include new packet forwarding information, or otherwise alter the data structure maintained by the sequence of pipelined on-chip memories.

At a step 212, the pre-lookup circuit 103 performs any pre-lookup operations. As noted above, the pre-lookup memory circuit 120, under control of the pre-lookup logic 115, performs a direct memory lookup in response to the lookup search key. As part of this step, the pre-lookup circuit 103 performs the following sub-steps:

- At a sub-step 212(a), the pre-lookup circuit 103 identifies the first d_0 bits of the lookup search key.

- At a sub-step 212(b), the pre-lookup circuit 103 performs a direct memory lookup to search for matching prefix values less than d_0 bits in length.
- At a sub-step 212(c), if the pre-lookup circuit 103 finds a matching prefix value, it records that prefix value in a pre-lookup best-match register (not shown).

As part of this sub-step, if the flag bits indicate that the lookup operation is to store information, the pre-lookup circuit 103 stores any new information, such as a new prefix value, in the identified location in the pre-lookup memory circuit 120.

At a step 213, the pre-lookup circuit 103 couples the lookup search key to the sequence of SRAM memories.

At a step 214, each stage of the sequence of SRAM memories, in parallel, performs a lookup on its portion of the lookup search key. As part of this step, the sequence of SRAM memories performs the following sub-steps:

- At a sub-step 214(a), each stage n in the sequence of SRAM memories identifies a span of d_n bits of the lookup search key. As noted above, the number of bits d_n might be different for each stage n in the sequence.
- At a sub-step 214(b), the n th on-chip SRAM memory performs a memory lookup for a portion of a prefix value, in response to at least a portion of the identified span of d_n bits.
- At a sub-step 214(c), the n th on-chip SRAM memory couples its lookup result to the corresponding n th address lookup circuit.

As part of this sub-step, if the flag bits indicate that the lookup operation is to store information, and the information to be stored is destined for a portion of the data structure in the n th on-chip SRAM memory, that memory stores any new information, such as a new prefix value, in the identified location therein.

5

In a preferred embodiment, as described in the incorporated disclosure, more than one stage of the on-chip SRAM memory can cooperate in response to a jointly-identified span of bits, to provide a single lookup result at the last one of those more than one stages.

10

Accordingly, as described in the incorporated disclosure, if the flag bits indicate that the lookup operation is to store information, more than one stage of the on-chip SRAM memory can cooperate in response to a jointly-identified span of bits, to store any new information, such as a new prefix value, in the identified locations in those stages of the on-chip SRAM memory.

15

At a step 215, the corresponding n th address lookup circuit processes the lookup result from the corresponding n th on-chip SRAM memory. As part of this step, the corresponding n th address lookup circuit determines if the lookup result is either a final lookup result (or a portion thereof), or an index into a corresponding stage $n+1$ in the sequence of SRAM memories. If the lookup result is a final lookup result (or a portion thereof), the corresponding n th address lookup circuit records the final lookup result (or portion thereof) in its corresponding n th best-match register (part of the corresponding n th pipeline register).

20

25

At step 216, the corresponding n th address lookup circuit couples its output to the corresponding stage n pipeline register. As noted above, the pipeline registers are clocked (preferably using a common clock) so that the lookup operation can proceed at each stage n substantially in parallel.

30

At a flow point 220, the method 200 is complete, and the final lookup result (if any) is coupled from the final stage best-match register to the output port 104.

5 As described above, as part of the method, if the flag bits indicate that the operation was to store (or otherwise modify or update) information in the data structure maintained by the pipelined memories, a result of the method 200 is to store information in the data structure synchronously with regard to lookup requests. Thus, each lookup request is performed with regard to a consistent data structure, because
10 storage operations occur explicitly either entirely before or entirely after the lookup request is performed, even though those lookup operations and storage operations are performed substantially concurrently by the pipelined memories.

Changes to the data structure that require more than one individual pass
15 through the sequence of pipelined memory stages are performed (first) by adding any necessary new entries to each memory, and then (second) by adding or altering a pointer that links the new entries to the main data structure accessible by lookup operations. Addition of the final pointer is an atomic act as viewed by lookup operations. Thus, the main data structure, as viewed by lookup operations, is always
20 consistent and is always updated atomically.

As described with regard to additional functionality above, it may occur that the final lookup result is coupled, along with the lookup search key, back to an earlier stage in the sequence of memories for further processing.

25

Generality of the Invention

The invention has general applicability to compact memory storage and retrieval, to update of information recorded in relatively compact form, and to
30 applications having utility for data lookup. None of these applications are limited specifically to lookup for message header information, nor are they necessarily

related to the specific applications disclosed herein. For example, embodiments of the invention can include one or more of, or some combination of, the following applications:

- 5 • Routers and switches at all levels of the ISO/OSI networking model, including without limitation: bridges, network gateways, network routers and switches.
- Emulation of an associative memory with contiguous prefix masking, including applications other than routing or switching.
- 10 • Both fixed-length and variable-length lookups in various types of lookup tables, such as for example for different protocols such as IP, Ethernet, MPLS, and different message types, such as for example unicast, multicast, and the like.

15

Moreover, techniques used by a preferred embodiment of the invention for lookup of message header information can be used in contexts other than the specific applications disclosed herein. For example, techniques used by embodiments of the invention for storage and retrieval of information in relatively compact form, and to relatively rapid pipelined data lookup, are all generally applicable to fields
20 other than the specific applications disclosed herein.

Other and further applications of the invention in its most general form would be clear to those skilled in the art after perusal of this application. The
25 invention would be usable for such other and further applications without undue experimentation or further invention.

Although preferred embodiments are disclosed herein, many variations are possible which remain within the concept, scope and spirit of the invention; these
30 variations would be clear to those skilled in the art after perusal of this application.

Claims

1. A method including concurrently performing a plurality of memory lookups in response to a sequence of inputs each having information;
5 each one of said memory lookups being performed in response to a corresponding distinct one of said inputs;
each one of said memory lookups being performed at a corresponding memory;
whereby each one of said sequence of inputs has a memory lookup
10 performed by an associated said memory for at least some of said information.
2. A method as in claim 1, including substantially concurrently providing results responsive to a plurality of said inputs accessing different subsequences of said memories.
- 15 3. A method as in claim 1, wherein at least some inputs have their information applied to all said memories at least once, and to at least some of said memories at least twice.
- 20 4. A method as in claim 1, wherein each one of said memory lookups is performed in response to a portion of said information.
5. A method as in claim 1, wherein each said input has substantially equal amounts of said information.
- 25 6. A method as in claim 1, wherein said memory lookups are each substantially performed on a single monolithic integrated circuit.
7. A method as in claim 1, wherein said sequence of inputs
30 includes at least one of: a destination IP address, an IP address, packet header information.

8. A method as in claim 1, wherein said corresponding memories collectively include lookup results including at least one datum responsive to each one of said inputs.

5

9. A method as in claim 8, wherein said lookup results collectively include packet forwarding information.

10. A method as in claim 1, wherein said memory lookup includes a sequence of individual memory accesses, each said individual memory access being performed at one of said memories.

11. A method as in claim 10, including substantially concurrently providing results responsive to a plurality of said inputs accessing different subsequences of said memories.

15

12. A method as in claim 10, wherein at least some inputs have their information applied to all said memories at least once, and to at least some of said memories at least twice.

20

13. A method as in claim 10, wherein each said input has substantially equal amounts of said information.

14. A method as in claim 10, wherein said sequence of individual memory accesses includes one said individual memory access at each said memory.

25

15. A method as in claim 10, wherein said sequence of individual memory accesses includes one said individual memory access at each said memory, followed by a second individual memory access at each said memory for at least a subsequence of said memories.

30

16. A method as in claim 10, wherein said sequence of individual memory accesses includes one said individual memory access at each said memory for only a subsequence of said memories.

5 17. A method including
coupling each one of a sequence of inputs to a sequence of memories, said sequence having a last memory and a next memory corresponding to each said memory other than said last memory, wherein each said memory is responsive to a distinct portion of said information;

10 coupling a result from each said memory other than said last memory to its corresponding said next memory in said sequence; and

providing an output of at least one of said memories;

whereby said sequence of inputs is each coupled to said sequence of memories in a pipelined manner to provide said output at a rate substantially equaling
15 one output as each input is received.

18. A method as in claim 17, including substantially concurrently providing results responsive to a plurality of said inputs accessing different subsequences of said memories.

20

19. A method as in claim 17, wherein at least some inputs have their information applied to all said memories at least once, and to at least some of said memories at least twice.

25 20. A method as in claim 17, wherein each said input has substantially equal amounts of said information.

21. A method as in claim 17, wherein said sequence of inputs includes at least one of: a destination IP address, an IP address, packet header
30 information.

22. A method as in claim 17, wherein said sequence of memories are substantially included in a single monolithic integrated circuit.

23. A method as in claim 17, wherein said output is responsive to a
5 sequence of individual memory accesses, each said individual memory access being performed at one of said memories.

24. A method as in claim 23, wherein at least some inputs have their information applied to all said memories at least once, and to at least some of said
10 memories at least twice.

25. A method as in claim 23, wherein said sequence of individual memory accesses includes one said individual memory access at each said memory.

15 26. A method as in claim 23, wherein said sequence of individual memory accesses includes one said individual memory access at each said memory, followed by a second individual memory access at each said memory for at least a subsequence of said memories.

20 27. A method as in claim 23, wherein said sequence of individual memory accesses includes one said individual memory access at each said memory for only a subsequence of said memories.

28. A method as in claim 17, wherein said sequence of memories
25 collectively include lookup results including at least one datum responsive to each one of said inputs.

29. A method as in claim 28, wherein said lookup results collectively include a set of packet forwarding information.

30

30. Apparatus including a circuit integrated onto a monolithic semiconductor chip, said circuit including

5 a sequence of registers each having a portion of a corresponding lookup search key in a sequence of said lookup search keys, each said register coupled to a corresponding one of a sequence of on-chip memories;

a subsequence of said memories not including a last said memory each having an output register associated therewith, said output register being coupled to an associated next said memory for each said memory in said subsequence;

10 a plurality of said memories capable of operating substantially concurrently each on a portion of a corresponding plurality of said lookup search keys;

whereby said sequence of registers is capable of coupling each lookup search key in portions to said sequence of memories, each said memory being responsive to each said lookup search key in sequence, each said lookup search key
15 being coupled to each said memory in sequence.

31. Apparatus as in claim 30, including an output register associated with said last memory and coupled to an output of said apparatus.

20 32. Apparatus as in claim 30, including an output register associated with said last memory, said output register being coupled to an associated earlier memory in said sequence.

25 33. Apparatus as in claim 30, wherein a plurality of output registers associated with different memories are each coupled to an output for said circuit.

34. Apparatus as in claim 30, wherein an output register associated with a memory other than said last memory is coupled to an output for said circuit.

35. Apparatus as in claim 30, wherein said lookup search key includes at least one of: a destination IP address, an IP address, packet header information.

5 36. Apparatus as in claim 30, wherein said memories collectively include lookup results including at least one datum responsive to each one of said inputs.

37. Apparatus as in claim 36, wherein said lookup results
10 collectively include packet forwarding information.

38. Apparatus including
a sequence of memories, said sequence having a last memory and a next memory corresponding to each said memory other than said last memory, each said
15 memory being coupled to a distinct portion of one of a sequence of lookup search keys;

each said memory other than said last memory being coupled to its corresponding said next memory in said sequence; and

whereby said sequence of inputs is each coupled to said sequence of
20 memories in a pipelined manner to provide said output at a rate substantially equaling one output as each input is received.

39. Apparatus as in claim 38, wherein a plurality of output registers associated with different memories are each coupled to an output for said circuit.
25

40. Apparatus as in claim 38, wherein an output register associated with a memory other than said last memory is coupled to an output for said circuit.

41. Apparatus as in claim 38, wherein said last memory is coupled to
30 an associated earlier memory in said sequence.

42. Apparatus as in claim 38, wherein said sequence of lookup search keys includes at least one of: a destination IP address, an IP address, packet header information.

5 43. Apparatus as in claim 38, wherein said sequence of memories collectively include lookup results including at least one datum responsive to each one of said lookup search keys.

10 44. Apparatus as in claim 43, wherein said lookup results collectively include a set of packet forwarding information.

1/2




Fig. 1

